A Comparison of Tomographic Reconstruction Methods for 3D Objects Using a Limited Number of Noisy Radiographic Views

by

*Thomas J. Asaki, Patrick R. Campbell, and Kevin R. Vixie

August 14, 2003

Abstract

This report provides visual and quantitative comparisons of 3D tomographic reconstructions from a limited number of radiographic views. Five 2D slices of representative single-material objects are reconstructed from nine noisy radiographic views using eight methods. The large variety of potential experimental situations does not allow us to make a complete study, but important conclusions can be made. We show that the quality of reconstructions depends significantly on the method employed and on the use of prior knowledge. In some cases, surprisingly good reconstructions are obtained when there is a good match between method, prior knowledge, and object morphology.

We employ the following basic inversion methods: projection matrix pseudo-inverse, filtered backprojection (FBP), and nonnegative least squares (NNLS). We also employ the following five noise-level constrained methods: truncated singular value decomposition (SVD), data-constrained total variation (TV), TV-prior SVD (TV-pSVD), nonnegative pSVD, and bounded pSVD.

The work reported here has been supported by the ASCI V&V Project and the Campaign 3 Simulation and Analysis Project under the direction of Allen Mathews of X-4.

*authors listed in alphabetical order

1 Introduction

Many methods exist for tomography using a few to several object views. Great reconstruction variablity exists due to the large indeterminacy of typical problems and the application of a variety of regularizations and prior knowledge constraints. We indtend to apply several methods to a set of standard test objects and compare the results both visually and quantitatively. Such a task cannot be exhaustive of all possible scenarios, but we will be able to qualitatively answer the following simple, yet important, question.

Given a certain class of objects and few radiographs, what reconstruction methods might best be used to provide certain pieces of information about the object?

We examined five test objects representing 2D slices of a 3D object. These objects have been described previously [2]. The objects are approximately represented on a 256 by 256 square voxel array. This is the $n = 65536$ dimensional object space. The reconstructions are performed on both a 256 by 256 voxel grid and on a 32 by 32 voxel grid. These are the two reconstruction spaces. They are thus of dimension 65536 and 1024, respectively.

Each object is projected onto a data space of nine radiographic views. The number of pixels per view is either 256 or 32 corresponding to the relevant reconstruction space. Thus the data spaces have dimension 2304 and 288, respectively. We have used an object mass projection so that the data represents an areal density, not a particle count. The data is given approximately Poisson additive noise of variable width dependent upon the size of the data space and the local data value. At 32 pixels per view the noise widths range from 1% for the minimal mass projection up to 5% at the maximal mass projection. At 256 pixels per view the noise widths range from 0.35% to 1.77%.

2 Essential Mathematics

In this section we provide the basic mathematics and notation necessary to discuss the methodologies and results. The descriptions are brief, but we provide a number of references throughout the text.

2.1 Projections

The data space and reconstruction space are connected by a linear projection operator. This projection matrix is completely determined by the geometry of the experiment. If the data is written as a vector d, the object mass distribution as a vector x, and the projection operator as the matrix P, then we have the simple relationship

$$d = Px. \tag{1}$$

The null space of P is very large and thus there is very significant ambiguity in the inversion of Eq. 1. Another way to think about it is that very many diverse objects project to exactly the same data. This is the central problem of sparse-data tomography. How does one choose the "correct" object from among many equally valid possibilites?

2.2 Singular Value Decomposition

If we apply singular value decomposition (SVD) to the matrix P we have that

$$d = U\Sigma V^T x \tag{2}$$

where U and V are orthogonal matrices whose columns span the data space D and reconstruction space X, respectively, and Σ is a diagonal matrix of nonnegative ordered singular values. This equation can be partially inverted by using a truncated set of singular vectors of U and V, provided that the corresponding singular values are nonzero. This SVD framework provides an initial working point for many reconstruction methods [1].

2.3 Prior-Based SVD

The use of prior information is easily incorporated into the SVD inversion framework [3]. If an object y is known to possess likely or definite object characteristics, then the null space projection of y can often be used as a correction term to an SVD reconstruction x. We have

$$x_{new} = x_{old} + (I - P)y. \tag{3}$$

If y is chosen on prior information and on properties of the current solution x_{old} then the process can be iterated – often with very good results.

2.4 Total Variation Regularization

Another essential procedure is total variation regularization (TV). The typical application is the denoising of images [12]. Suppose we have an initial noisy image $\hat{x}$. The TV-regularized image x is the solution to

$$min\{\mathrm{TV}(x) + \alpha \parallel x - \hat{x} \parallel^2\}_{x \in X}, \tag{4}$$

where $\parallel . \parallel$ indicates the L^2 norm and the TV functional applied to a continuous function is

$$\mathrm{TV}(f(t)) = \int_\Omega |\vec{\nabla} f(t)| dt \tag{5}$$

The TV functional for discrete distributions is a matter of interpretation of Eq. 5, but is most often computed using standard forward and backward differencing schemes.

The application to tomography is relatively new [4, 5]. The choice among all possible reconstructions x that satisfy the data d is given by

$$min\{\mathrm{TV}(x) + \alpha \parallel Px - d \parallel^2\}_{x \in X}. \tag{6}$$

2.5 Filtered Backprojection

The filtered backprojection (FBP) method of tomography is the standard in the field of medical imaging and for other applications for which a large number of radiographs are available [9]. The basic concept is to perform a Fourier transform in the data space, apply

an appropriate filter (for denoising and data weighting), and sum the individual inverse transforms in the reconstruction space.

2.6 Metrics

The quality of a reconstruction depends entirely upon the experimenter's goals. Some possible metrics are to determine within some tolerance: object placement, total mass, object density at specified points, material identifications, and boundary shapes. Clearly, the list is endless. For our illustrative purpose, the quality of a reconstruction is given by the L^2 norm over the all densities in the reconstruction space:

$$\frac{\| x - x^* \|^2}{\| x^* \|^2}, \tag{7}$$

where x^* is the true object and x is a 256 by 256 rescaled version of a reconstruction.

3 Methods

We describe here some of the mechanics of the reconstruction methods and discuss general results pertaining to the five test objects. The descriptions are intended to provide an adequate basis for comparisons for someone familiar with tomographic reconstruction. Further descriptions can be found in the following references: SVD [8, 6, 1], pSVD [3], TV [12], FBP [9], and NNLS [10, 12].

Figures 1 through 5 show 32 by 32 reconstruction slices of five standard test objects from nine simulated noisy radiographs. The false color scheme shows the object and reconstruction mass densities at each voxel location. The nine radiographs are computed for equally-spaced angles from zero to π radians and for a particle beam in the plane of the 2D object slice. The norm (Eq. 7) for each reconstruction is shown in the lower left corner of each subfigure. Figures 6 through 10 show reconstructions of the same objects using the same methods except the detector resolution and reconstruction resolution are eight times finer. The simulated noise level for the detectors is reduced by a factor of $\sqrt{256/32} = \sqrt{8}$.

Object. The objects are normalized so that their densities lie between zero and unity. The colorbar scale is the same for each figure. Object #3 has exact densities of $(0, 0.25, 0.50, 0.75, \text{and } 1.00)$ providing convenient visual comparison colors for all other figures.

Projection Matrix Inversion. The projection matrix is the linear operator that projects the object-space onto the data-space. The projection is inverted using SVD retaining all singular vectors corresponding to nonzero singular values. This fullest possible inversion (FPI) attains a best possible fit to noisy data at the expense of possible spurious features in the reconstructed object. The full inversion reconstructions at a resolution of 32 are so noisy as to bear very little resemblance to the original object. The reconstructions at a resolution of 256 are better due to the reduced noise level chosen for the data. Even so, the reconstructions capture only the crudest detail.

Filtered Backprojection. The filtered backprojection solution (FBP) is obtained using the publically available Matlab code of Adel Faridani [7]. A ramp filter was used and the result was mass corrected. FBP provides an expected Fourier-like smoothed reconstruction. Sharp boundaries in the object show up smoothed and there are underlying 18-fold symmetry artifacts that are a direct result of the nine view input. The density values at each extreme saturate the color scale.

Nonnegative Least Squares Solution. The nonnegative least squares solution (NNLS) is the reconstruction x obtained by the Matlab function calls

```
ops=optimset('TolX',.0001);
x=lsqnonneg(P,d,[],ops);
```

The algorithm [10] starts with a set of possible basis vectors and computes the associated dual vector λ. It then selects the basis vector corresponding to the maximum value in λ in order to swap out of the basis in exchange for another possible candidate. This continues until $\lambda \leq 0$. The NNLS reconstruction does a good job with the preservation of the zero density background (outside of the object). Positive density regions of the object are poorly determined – again because the solver attempts to accurately match noisy data. Many of the high-density values saturate the color scale. NNLS reconstructions are not given at 256 by 256 resolution because of the overwhelming time required by the algorithm to run to completion ($\sim$ years).

Truncated SVD Inversion. The projection matrix is partially inverted using a truncated set of singular vectors corresponding to the singular value set

$$\sigma_i \leq \sigma_1 \delta \tag{8}$$

where δ is a measure of the noise amplitude across the data. For these tests we take

$$\delta = \| \frac{\epsilon_{hi}^{f} \epsilon_{lo}^{1-f}}{2ln(\epsilon_{hi}/\epsilon_{lo})} \| \frac{1}{\| f \|} \tag{9}$$

according to the artificial noise ϵ added to the data (see Introduction) and the relative mass projection $f = d/d_{max}$. In this way we reconstruct object features without noise amplification. This noise-level constrained inversion is the best low-dimensional reconstruction in the absence of prior knowledge and regularization schemes [1]. This reconstruction is naturally noise-supressing, but does suffer from $2v$-fold reconstruction artifacts. This reconstruction also provides a Fourier-like smoothed version of the object – high spatial frequency features are lost.

Total Variation Regularized. The reconstruction x is the solution of Eq. 4 iteratively solved for a suitable α such that $\| Px - d \|^2 \approx \delta^2$. Each minimization is solved by the fixed-point method and the TV is calculated using a forward and backward differencing scheme. The Matlab code is a slightly modified version of Kurt Vogel's publically available code [11]. The TV-regularized reconstruction allows for accurate reconstruction of both high- and low-spatial frequency features as long as the projection

data is satisfied to within a noise constraint. This solution is particularly useful when attempting to reconstruct density discontinuities and constant density regions. Thus, the zero density regions are well captured as are sharp object-boundaries. TV does have the drawback that density levels are changed; so while the visual effect is impressive, the density values may be altered. However, the noise level constraint helps to minimize this effect.

pSVD with TV Prior. The TV prior is the object y which satisfies the minimization requirement Eq. 4 using an α approximately three times larger than that determined for the TV reconstruction described above. The process is iterated according to Eq. 3 until $\| Px - d \|^2 \approx \delta^2$. Different choices of α can lead to different final results. A large α will tend the solution to a reconstruction similar to the truncated SVD. A small α will tend the solution to a reconstruction similar to the TV.

pSVD with Nonnegativity Prior. The nonnegativity prior is the object y satisfying $y = max(0, x)$. The process is iterated according to Eq. 3 until $\| Px - d \|^2 \approx \delta^2$. This reconstruction has the advantange of nudging the solution toward a data-satisfying positive-density object. Typically, the solution is desirable over the truncated SVD solution because it reduces $2v$-fold symmetry artifacts and more accurately identifies sharp density discontinuities [3].

pSVD with Bounded Density Prior. The bounded density prior is the object y satisfying $y = min(1, max(0, x))$. The process is iterated according to Eq. 3 until $\| Px - d \|^2 \approx \delta^2$. This reconstruction method is identical to that previous except that the density prior is bounded above and below. It works exremely well for cases in which the limiting densities are well known and comprise a large part of the object (see Objects #1 and #2).

4 Discussion

We present here observations on the comparisons between the several reconstructions.

1. The fullest possible projection matrix inversions (FPI) are noisy beyond much usefulness. They are shown here mainly for illustrative purposes; constraints and regularizations are essential for the interpretation of *noisy* data.
2. The FBP reconstructions exemplify the well-known result that too few views are inadequate. The reconstructions contain significant artifacts associated with the detector placements and some large density deviations. FBP performs best in the central region of the reconstruction.
3. The NNLS solutions are all quite poor because they attempt to exactly fit the noisy data. The one merit of the method is that it does very well at capturing the location of zero-density regions. Unfortunately, it also adds many zero-density regions where none exist.
4. The truncated SVD reconstructions (trSVD) are our first attempt at a constrained reconstruction based upon our knowledge of the noise characteristics of the data. The

overall performance is much better than the basic methods (FBP, FPI, and NNLS). Better solutions than trSVD must include some regularization.

5. The visual quality of the trSVD reconstructions at 32 by 32 resolution is similar to that of FBP. But note that no particular region of the reconstruction space is favored, the reconstruction projected data has the right characteristics and high amplitude deviations are absent. However, much of the $2v$-fold reconstruction artifacts are the same.
6. The visual quality of the trSVD reconstructions at 256 by 256 resolution is superior to that of FBP. For the more complex objects the contrast is poorer but the densities are much better matched across the whole reconstruction space.
7. The σ-constrained TV reconstructions are our first attempt at providing a regularizing condition. Since TV tends to eliminate voxel-size features while preserving edges and smoothness, we find that reconstruction noise is greatly reduced, constant density regions are well reproduced and there is no noise concentration effects seen in many pSVD methods. However, there are noticeable changes in narrow or small features such as the crescent tips of object #3 and the high spatial-frequency waves of object #4.
8. The nonnegativity pSVD solution is fundamentally different than the NNLS solution. The zero-level preservation is similar to that of TV but at the expense of concentrating the noise into the nonzero density regions. This is especially evident when examining the first three objects.
9. Of the three pSVD methods used, all show significant enhancement of sharp density boundaries. The reconstruction merit is similar for all.
10. The TV-pSVD appears to have the greatest potential as the best reconstructor. Its ability to capture a particular object however, depends upon one's choice of α. The same choice of α was made for all figures yet the Fig. 6 reconstruction is superior while the Fig. 8 is not clearly the best.
11. Not surprisingly, the bounded-density prior works exrememly well for objects of binary and known densities. Consider the reconstructions of Figs. 2 and 3 remembering that the data is noisy and we have only nine views. A true binary density reconstruction is even more impressive (see many examples in Ref. [3]).
12. Both TV and bounded-density pSVD methods are notable for their ability to suppress spurious density values.
13. All reconstructions can be expected to improve with more views and less noise. This exploration is beyond the scope of this report.

References

[1] T.J. Asaki and K.R. Vixie. SVD analysis for radiographic object reconstruction I: Initial results. Technical Report LA-UR-01-6534, Los Alamos National Laboratory, 2001.

[2] T.J. Asaki and K.R. Vixie. Standard 2D test objects for radiographic inversion studies. Technical Report LA-UR-02-3978, Los Alamos National Laboratory, 2002.

[3] T.J. Asaki and K.R. Vixie. SVD analysis for radiographic object reconstruction II: Null space enhancements driven by minimal constraints. Technical Report LA-UR-pending, Los Alamos National Laboratory, 2003.

[4] T.J. Asaki, K.R. Vixie, and P.R. Campbell. Total variation regularization applied to tomographic reconstruction. Technical Report (in preparation), Los Alamos National Laboratory, 2003.

[5] Paul Channell. Consistency of multiple view radiographic reconstructions. Technical Report LANSCE-1:03-012(TN), Los Alamos National Laboratory, 2003.

[6] James W. Demmel. *Applied Numerical Linear Algebra.* SIAM, 1997.

[7] Adel Faridani. *http://www.owlnet.rice.edu/ elec431/projects96/DSP/index.html.*

[8] Gene H. Golub and Charles F. VanLoan. *Matrix Computations.* Johns Hopkins University Press, 1996.

[9] Avinash C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging.* SIAM, 2001.

[10] C.L. Lawson and R.J. Hanson. *Solving Least Squares Problems.* Prentice-Hall, 1974.

[11] Curtis R. Vogel. *http://www.math.montana.edu/ vogel/Book/Codes/Ch8/2d/.*

[12] Curtis R. Vogel. *Computational Methods for Inverse Problems.* SIAM, 2002.

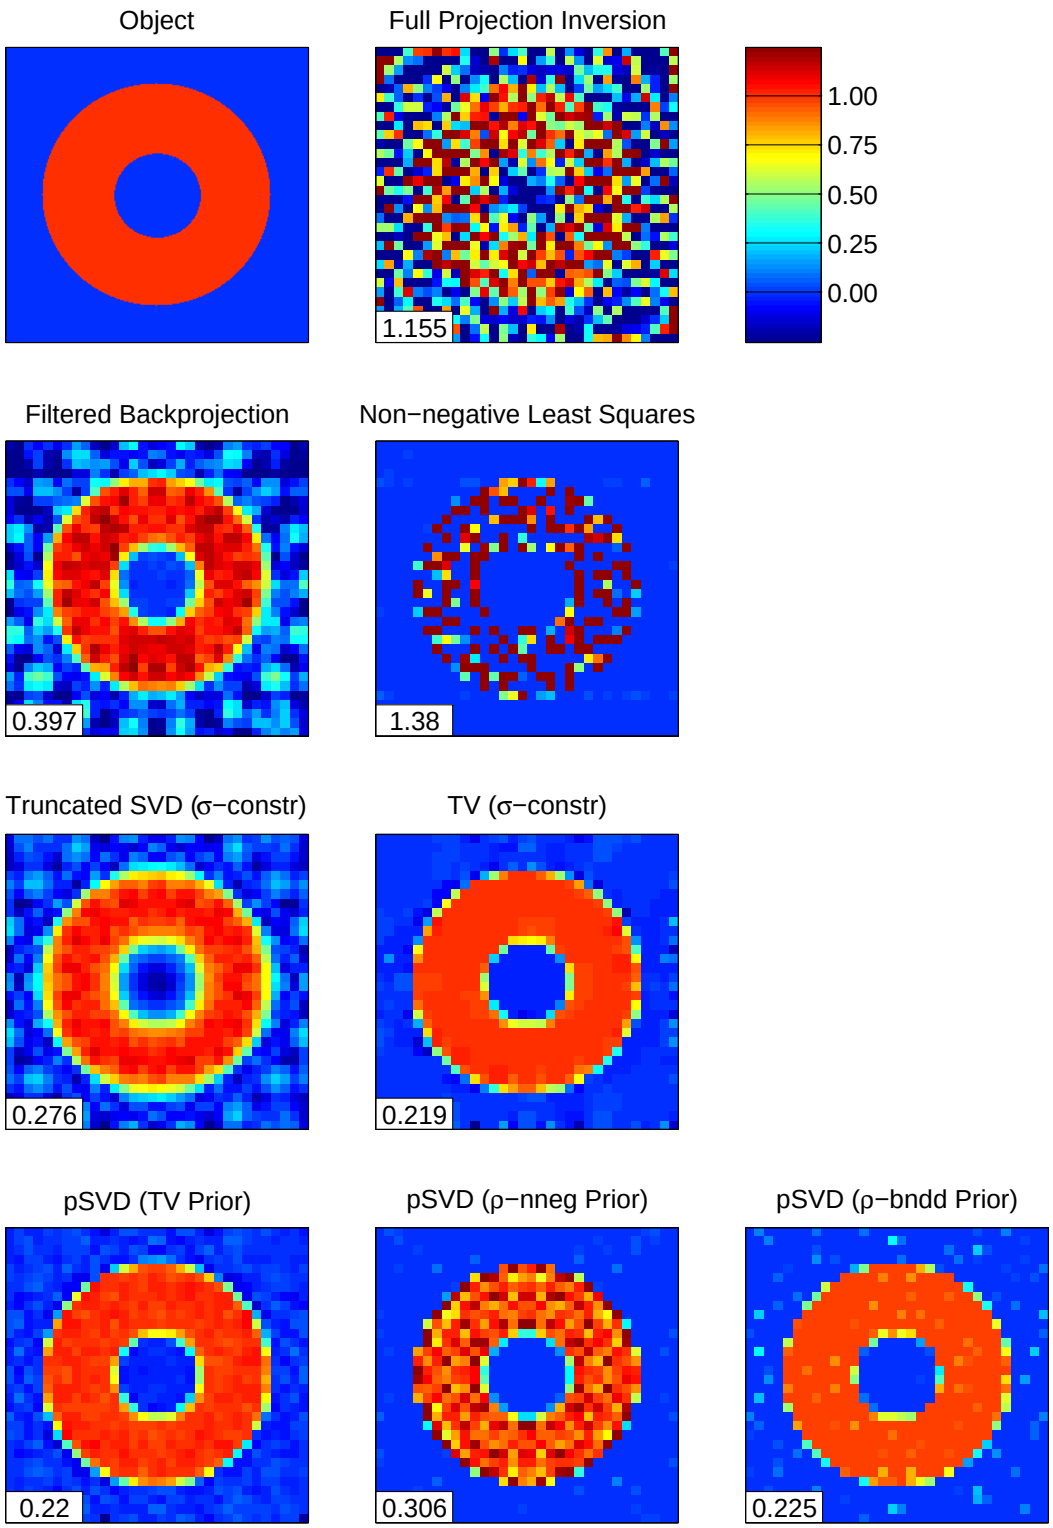

Figure 1: Object #1 (256 by 256) and several nine-view reconstructions (32 by 32). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

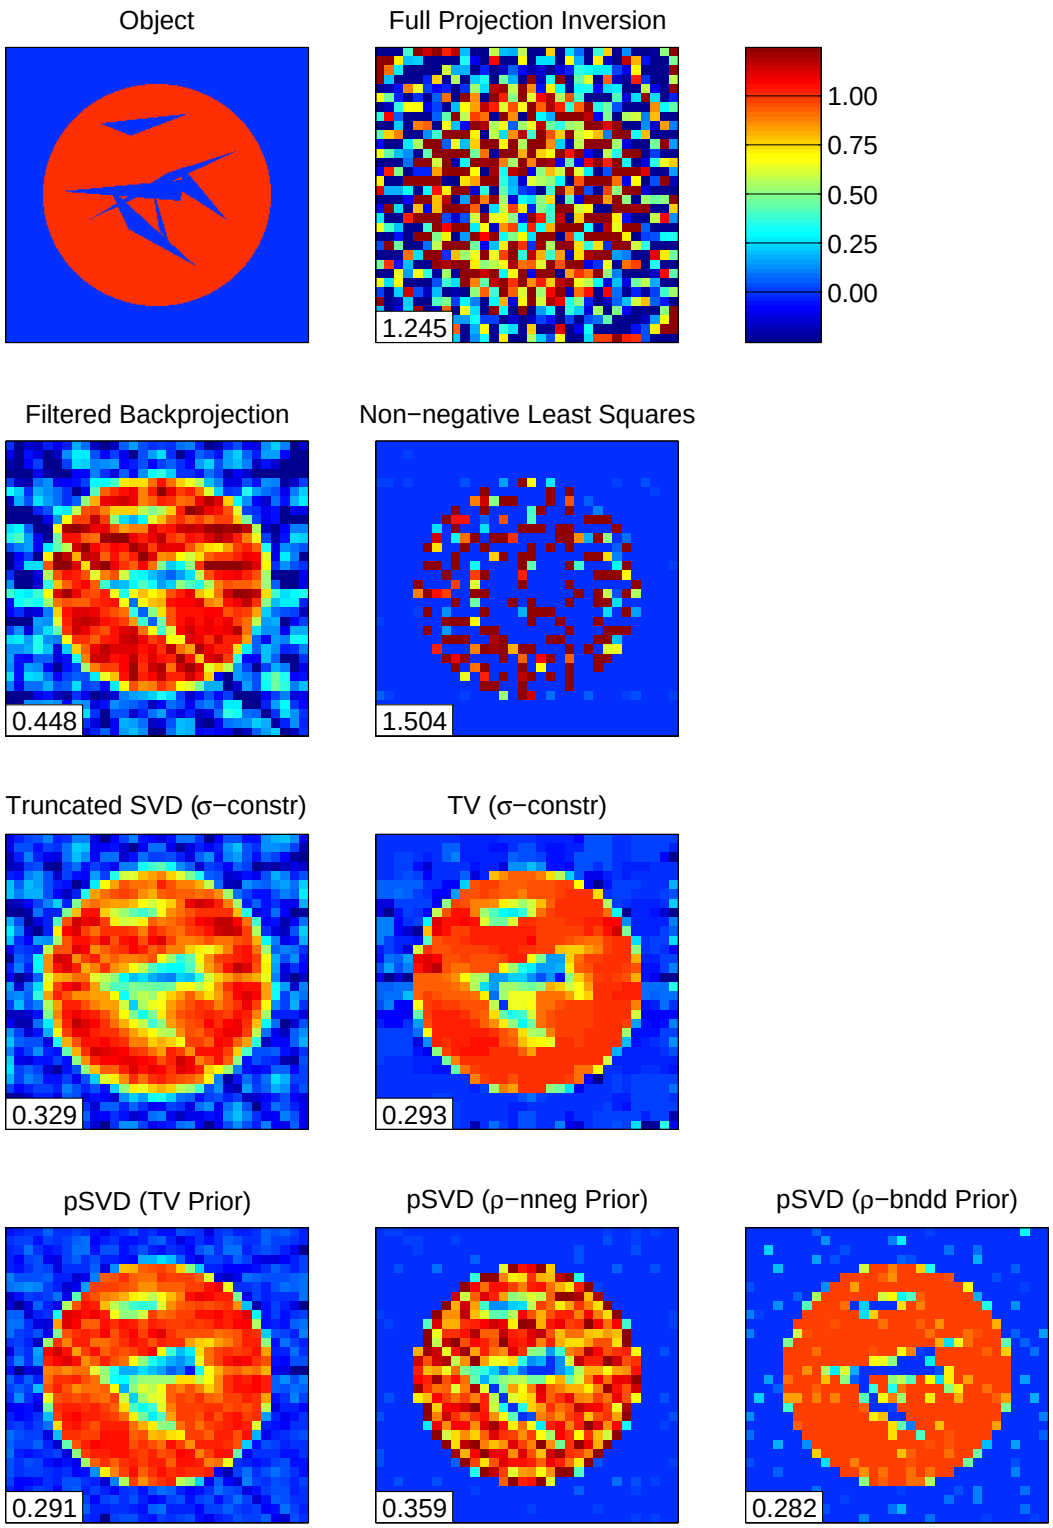

Figure 2: Object #2 (256 by 256) and several nine-view reconstructions (32 by 32). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

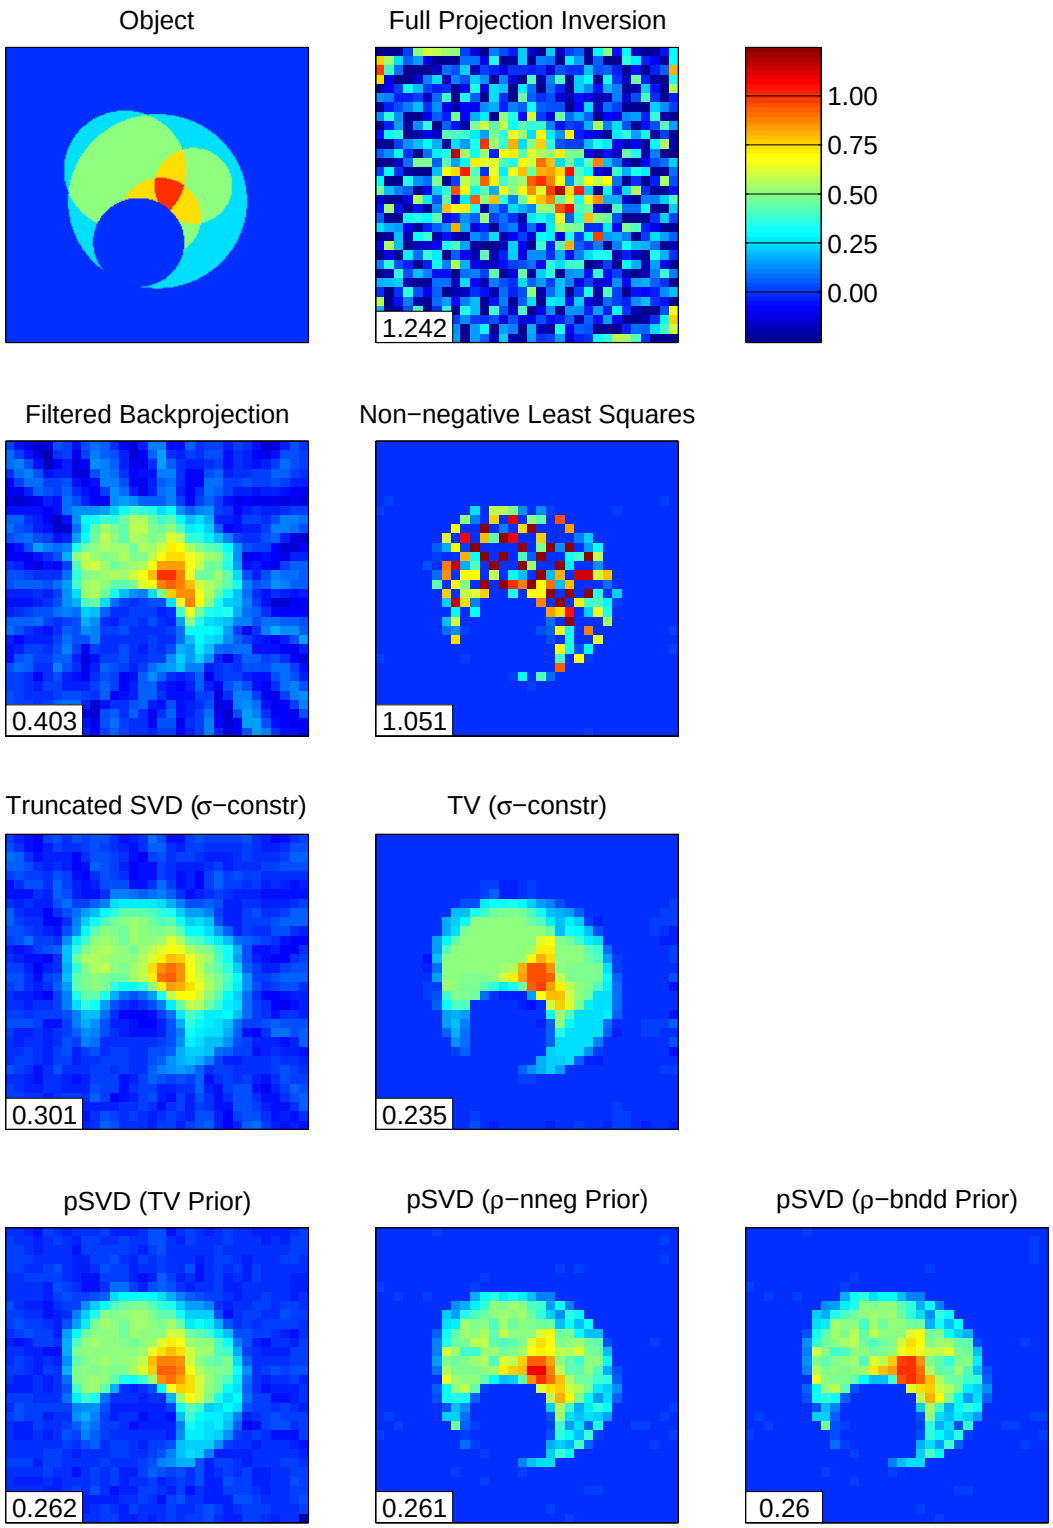

Figure 3: Object #3 (256 by 256) and several nine-view reconstructions (32 by 32). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

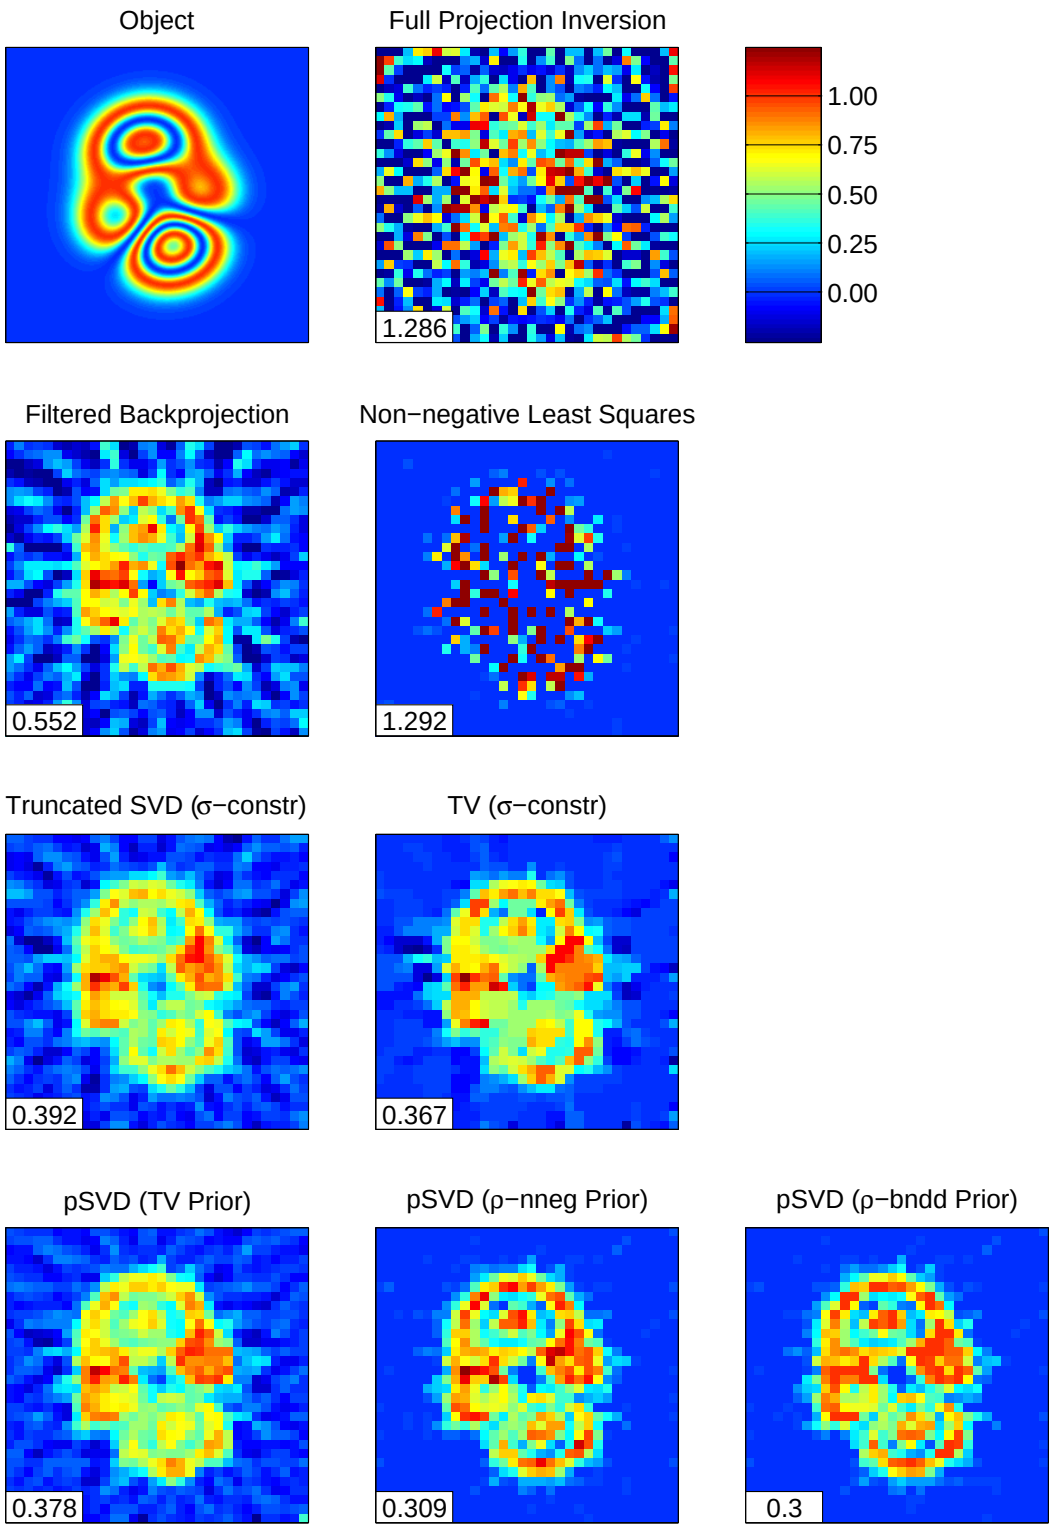

Figure 4: Object #4 (256 by 256) and several nine-view reconstructions (32 by 32). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

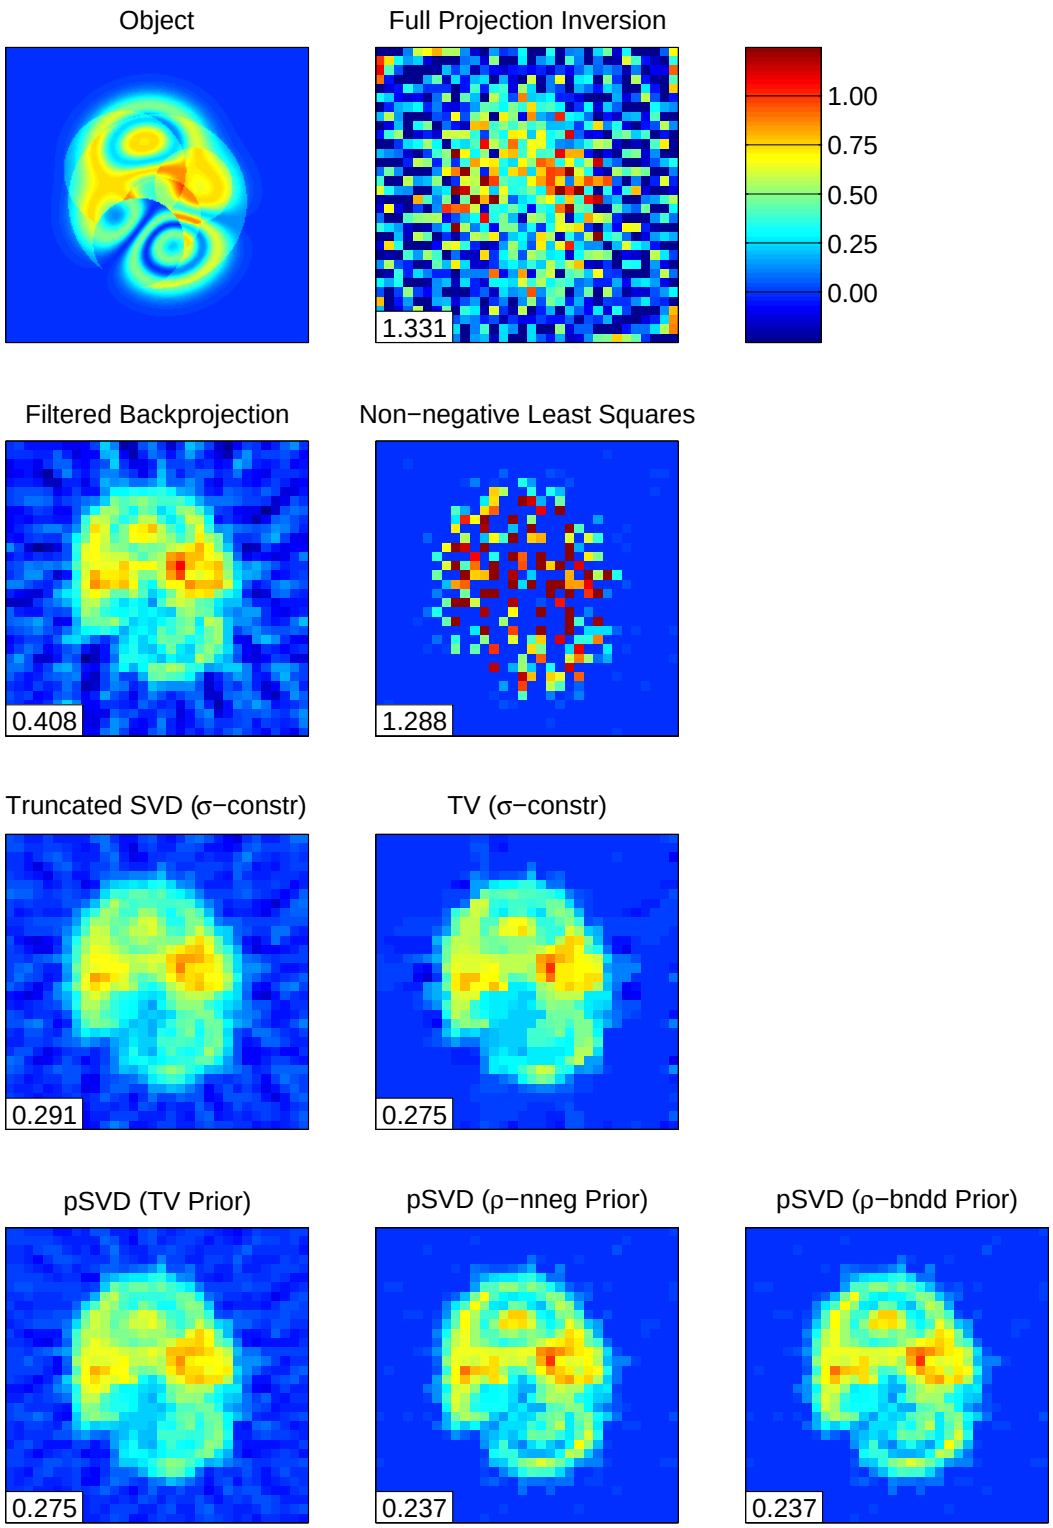

Figure 5: Object #5 (256 by 256) and several nine-view reconstructions (32 by 32). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

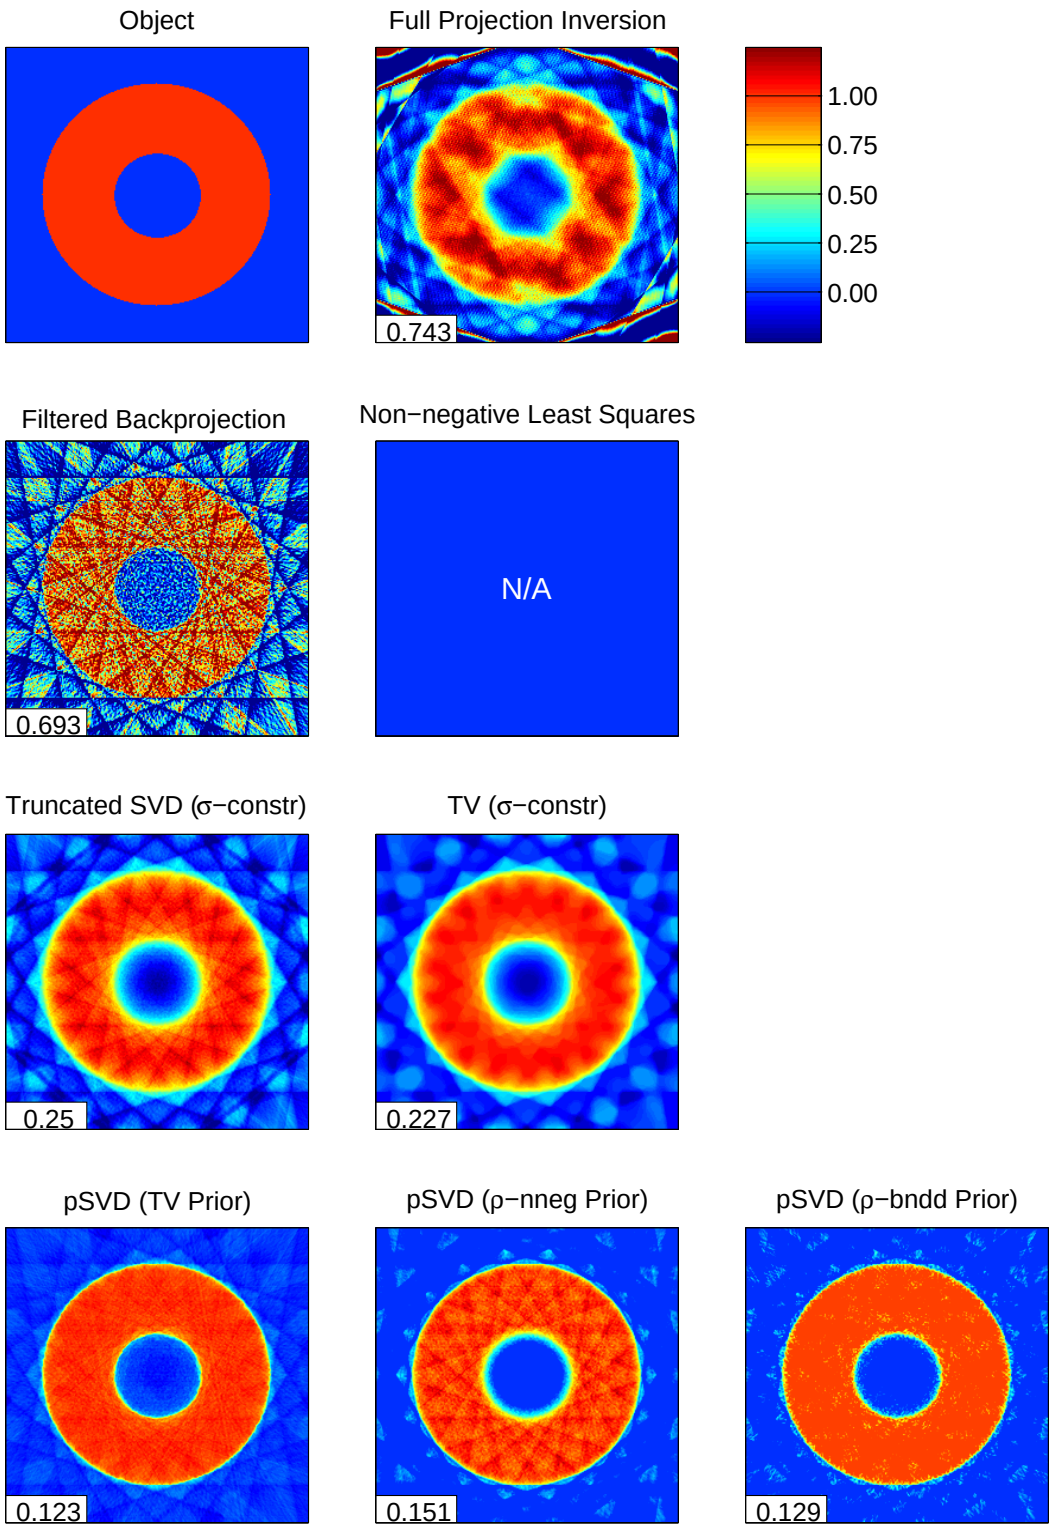

Figure 6: Object #1 (256 by 256) and several nine-view reconstructions (256 by 256). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

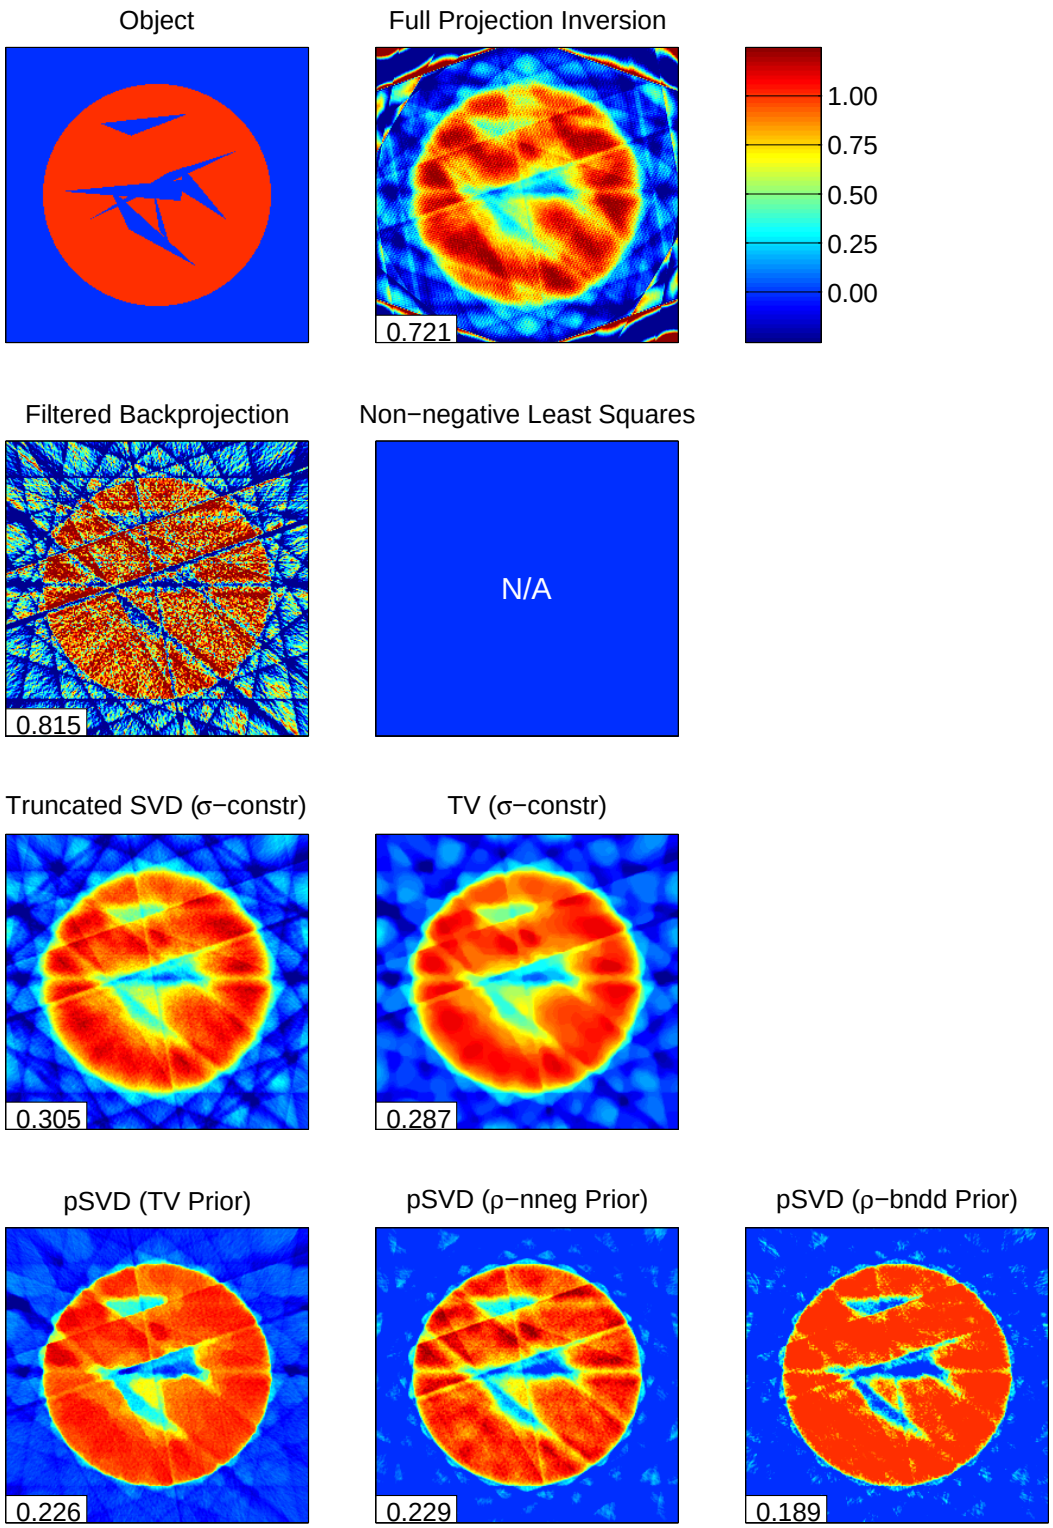

Figure 7: Object #2 (256 by 256) and several nine-view reconstructions (256 by 256). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

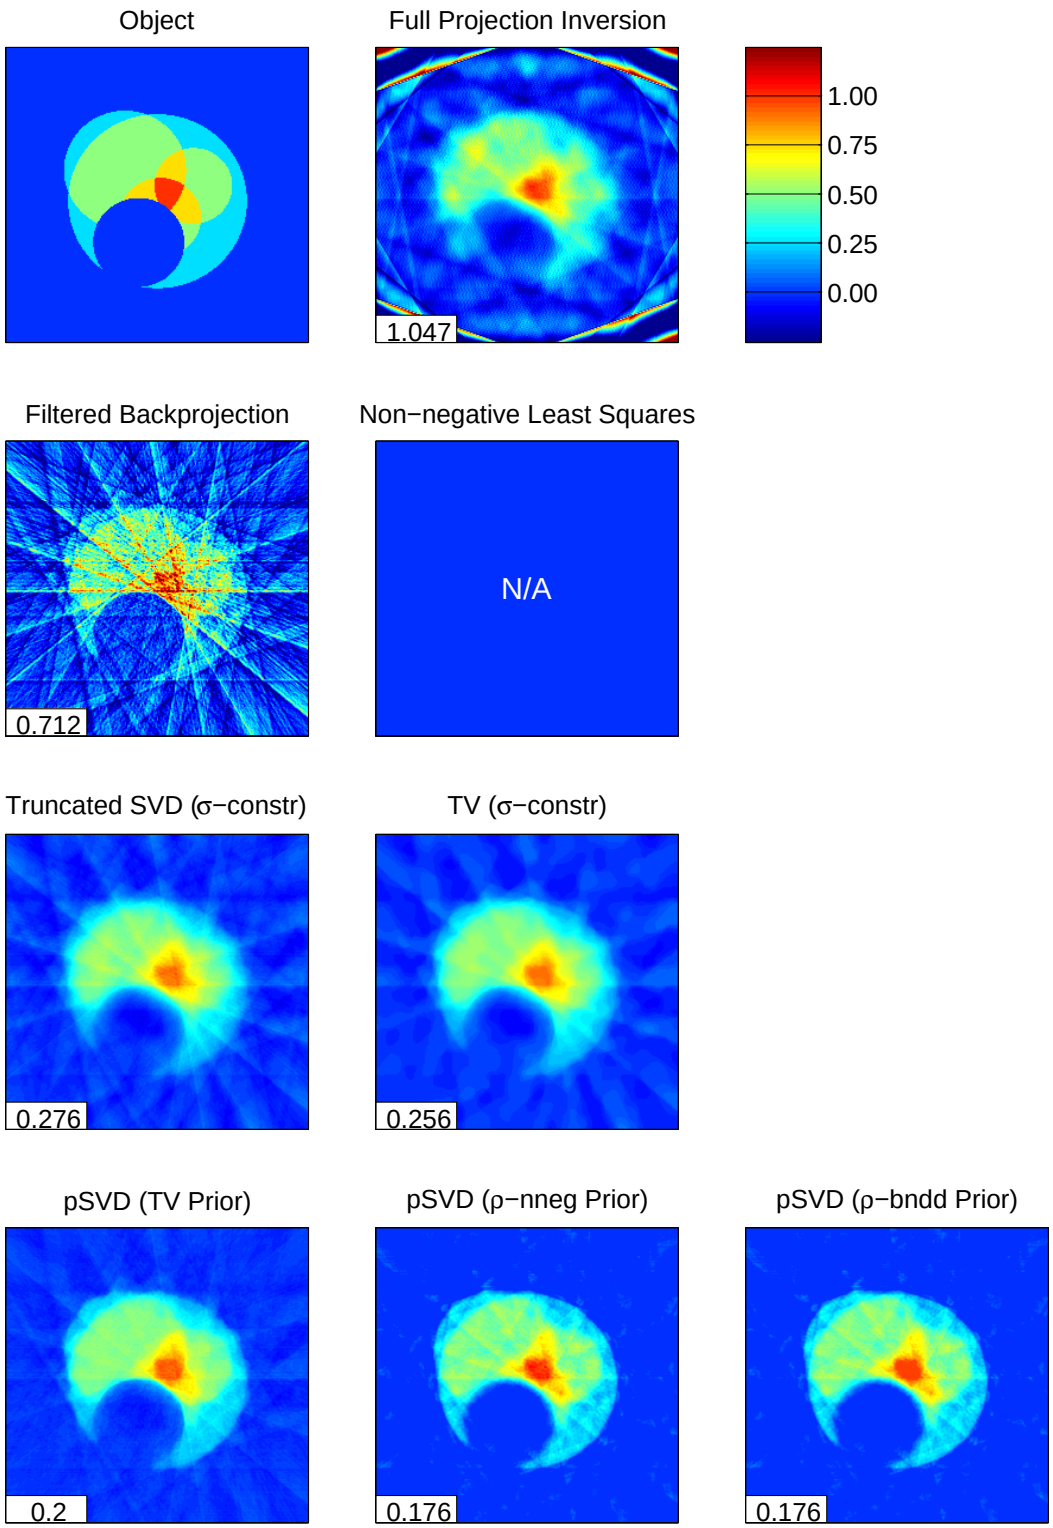

Figure 8: Object #3 (256 by 256) and several nine-view reconstructions (256 by 256). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

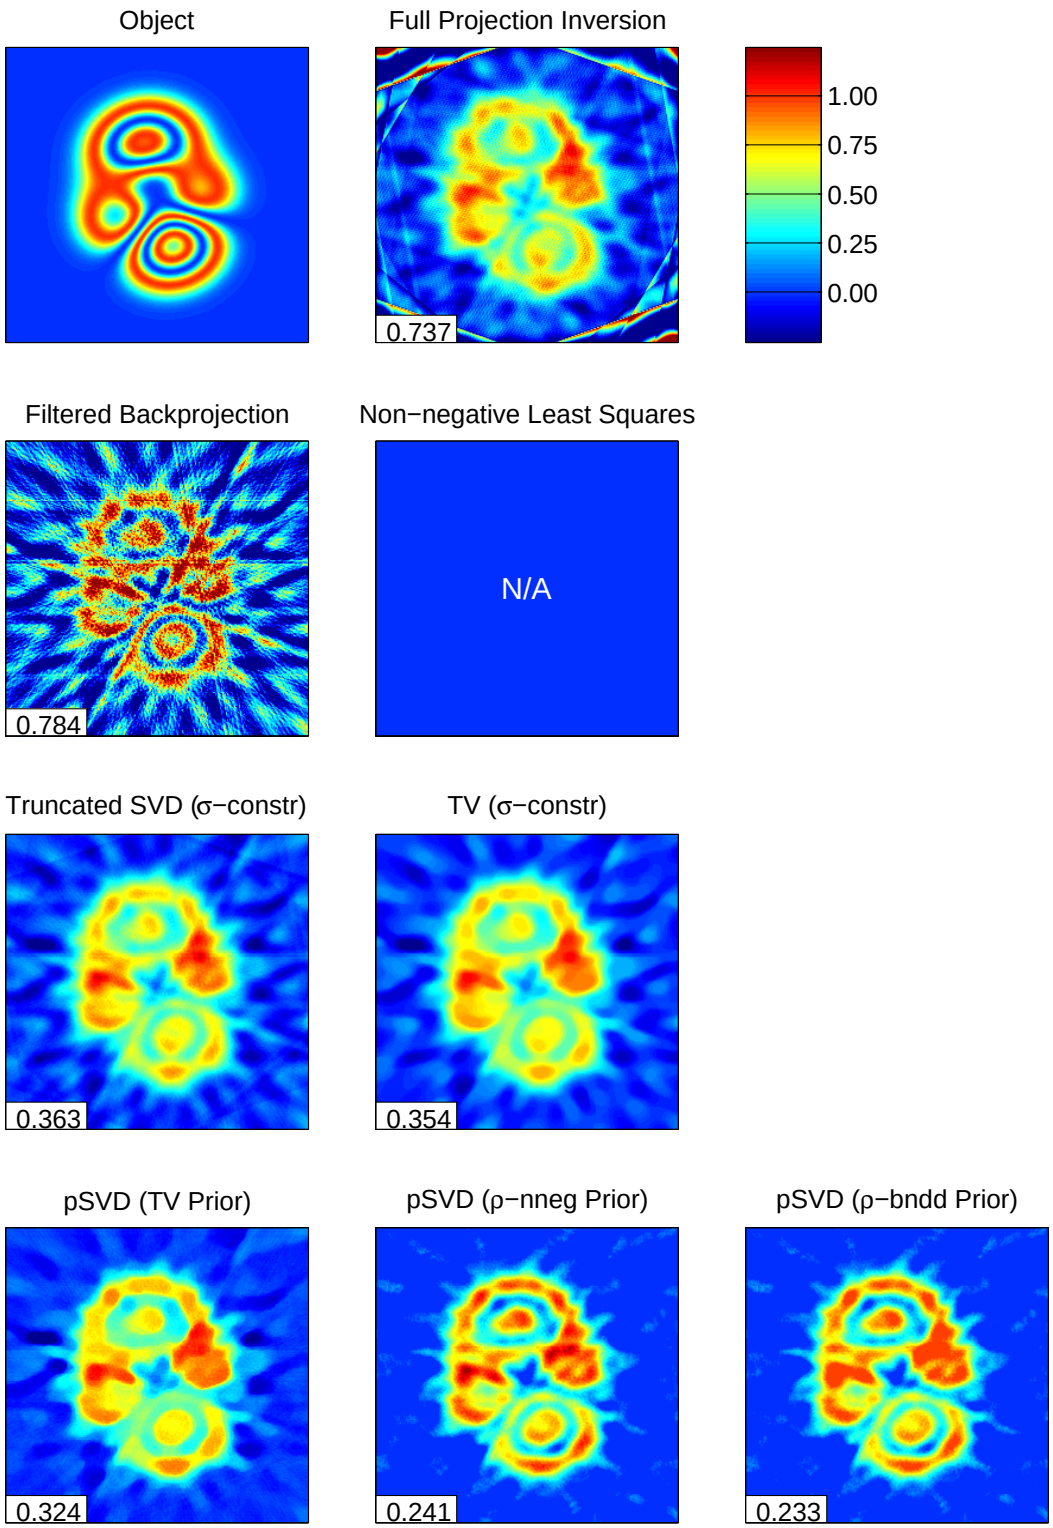

Figure 9: Object #4 (256 by 256) and several nine-view reconstructions (256 by 256). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.

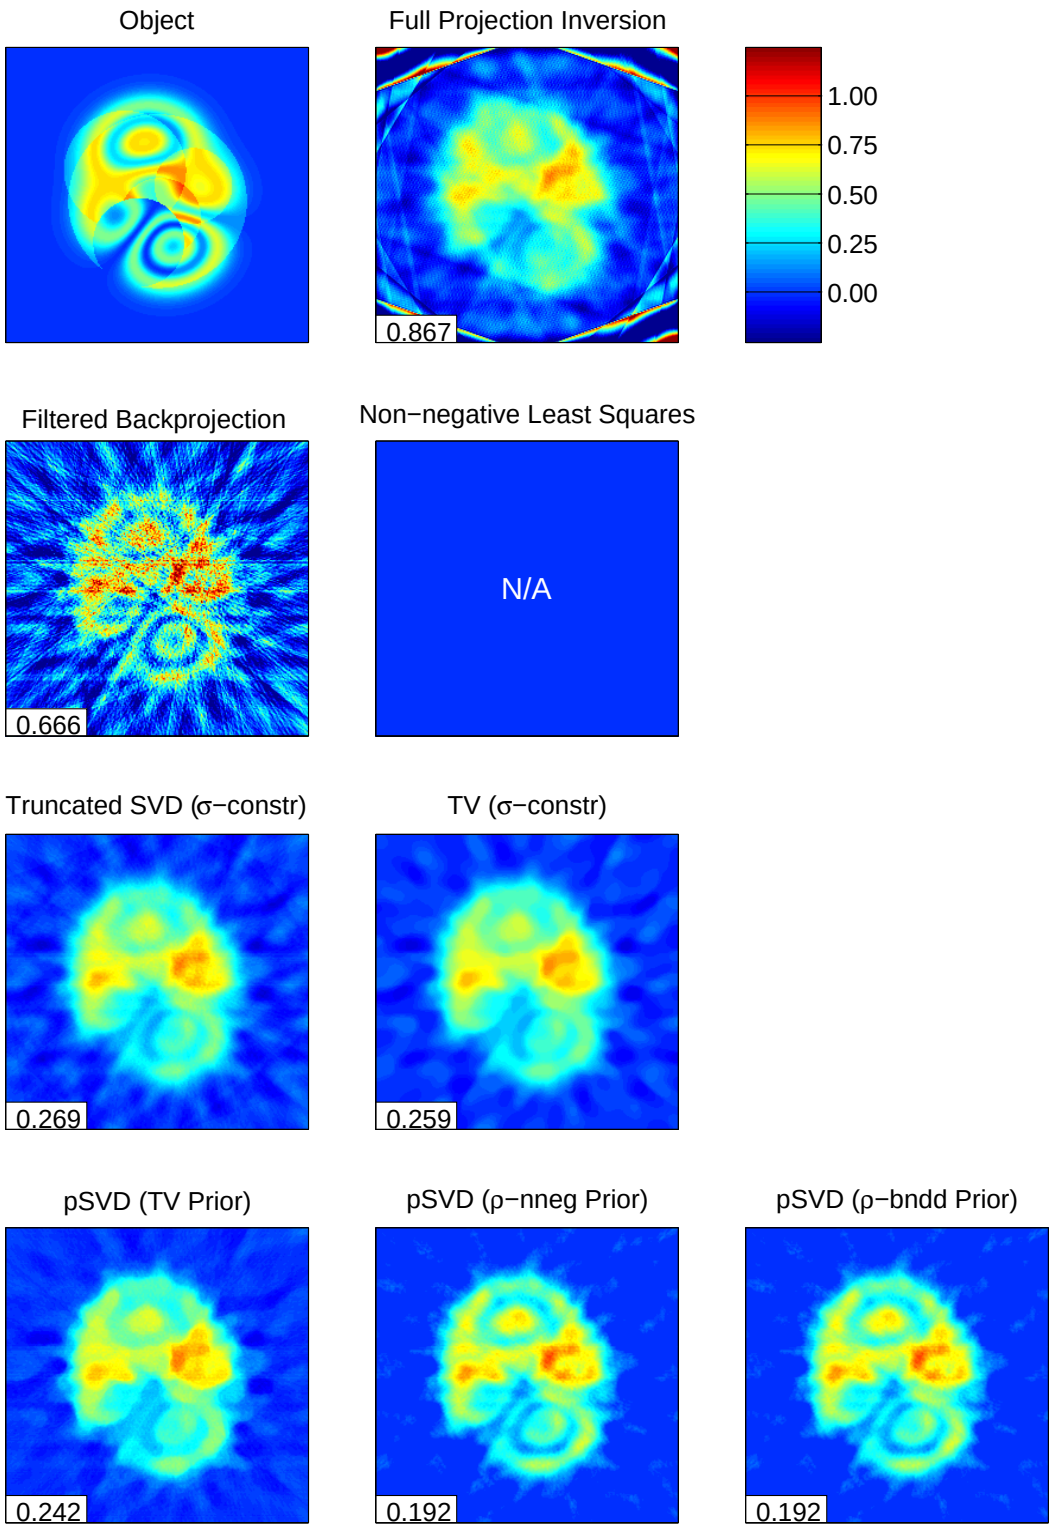

Figure 10: Object #5 (256 by 256) and several nine-view reconstructions (256 by 256). L^2 norms are shown in the lower left of each reconstruction. Method details are given in the main text.